

HandDAGT: A Denoising Adaptive Graph Transformer for 3D Hand Pose Estimation

Wencan Cheng¹[0000-0002-7996-0236], Eunji Kim¹[0009-0005-6039-1853], and Jong Hwan Ko²[0000-0003-4434-4318]

¹ Department of Artificial Intelligence, Sungkyunkwan University, South Korea

² College of Information and Communication Engineering, Sungkyunkwan University, South Korea

{cwc1260, dmsw1dpvm, jhko}@skku.edu

Abstract. The extraction of keypoint positions from input hand frames, known as 3D hand pose estimation, is crucial for various human-computer interaction applications. However, current approaches often struggle with the dynamic nature of self-occlusion of hands and intra-occlusion with interacting objects. To address this challenge, this paper proposes the Denoising Adaptive Graph Transformer, HandDAGT, for hand pose estimation. The proposed HandDAGT leverages a transformer structure to thoroughly explore effective geometric features from input patches. Additionally, it incorporates a novel attention mechanism to adaptively weigh the contribution of kinematic correspondence and local geometric features for the estimation of specific keypoints. This attribute enables the model to adaptively employ kinematic and local information based on the occlusion situation, enhancing its robustness and accuracy. Furthermore, we introduce a novel denoising training strategy aimed at improving the model’s robust performance in the face of occlusion challenges. Experimental results show that the proposed model significantly outperforms the existing methods on four challenging hand pose benchmark datasets. Codes and pre-trained models are publicly available at <https://github.com/cwc1260/HandDAGT>.

Keywords: Hand pose estimation · Point cloud · Transformer

1 Introduction

3D hand pose estimation has emerged as a pivotal element in various human-computer interaction systems, providing essential keypoints for human hands. Accurate localization of hand keypoints enables precise understanding and effective analysis of human motion and behavior. With the advancements in deep learning techniques and the proliferation of low-cost depth cameras, significant progress has been made in depth image-based 3D hand pose estimation.

Early deep learning-based methods showcased the efficacy of conventional convolutional neural networks (CNNs) in hand pose estimation tasks [4, 10, 11, 14, 17, 36, 38, 41]. Subsequently, CNN architectures leveraging different modalities,




Fig. 1: Illustration of the HandDAGT concept. The model feeds 3D local patches cropped from input depth images and corresponding point clouds to the adaptive graph transformer for the keypoint coordinate estimation. The local patches are disturbed by the random Gaussian noise during training for a robust performance.

such as point clouds [6,8,13,16], voxelized representations [15,31], or multi-modal combination [35], have been extensively explored. Beyond CNN, the transformer-based approaches [21,45] also showed the feasibility in the hand pose estimation field. Despite the notable successes achieved by these approaches, achieving accurate estimation remains challenging in scenarios such as self-occlusion and hand-object occlusion. Recent studies [6,35,45], have endeavored to introduce a static graph structure to capture kinematic correspondence between keypoints, enabling the reconstruction of occluded keypoints using corresponding keypoints’ information. However, static graph structures fall short of capturing the dynamics of various occlusions.

To address these challenges, we propose a novel denoising adaptive graph transformer, named HandDAGT, to adaptively weigh local detail and kinematic correspondence for accurate hand pose estimation across various occlusion conditions. The proposed transformer can learn precise local geometric information of the hand as the traditional local transformers [12,46,49]. Importantly, the transformer is meticulously designed to calculate the geometric feature and kinematic topology information directly in the 3D keypoint space. This design is made to circumvent 2D-to-3D non-linearity that arises when operating in 2D space [6]. Moreover, for accurate estimation, it integrates a novel adaptive attention module that adaptively activates local geometric information or information from kinematically corresponding keypoints depending on whether keypoints are well-presented or occluded. Furthermore, to enhance the robustness and accuracy of estimation under occlusion, we introduce a novel denoising training strategy as it compels the model to correct estimations even with disturbed patches.

We evaluate HandDAGT on four challenging benchmarks, including single-hand ICVL [40] and NYU [41] dataset, and hand-object DexYCB [2] and HO3D [18] datasets. The results show that our network achieves the state-of-the-art performance with the lowest mean distance errors of 5.66 mm and 7.12 mm on the single-hand ICVL and NYU datasets, respectively. The model also achieves state-of-the-art performance on the hand-object DexYCB and HO3D datasets with the lowest error of 8.03 mm and 1.81 cm, respectively.

Our main contributions are summarised as follows:

- We propose a novel transformer-based architecture, HandDAGT, for accurate 3D hand pose estimation that utilizes the depth image and point cloud as a multi-modal input.
- We propose a novel adaptive attention module designed to weigh the local geometric information and kinematically evolved keypoint feature for high reliability to occlusions.
- We introduce a novel denoising training strategy for hand pose estimation to improve robustness and accuracy.
- We perform comprehensive experiments on big and challenging benchmarks that present the new state-of-the-art performance of our proposed method.

2 Related Work

2.1 3D Hand Pose Estimation Based on Depth Image.

In the realm of 3D hand pose estimation utilizing depth images, traditional 2D CNN-based approaches [4, 10, 11, 14, 17, 36, 38, 41] have garnered widespread adoption due to their straightforward implementation. However, they are encumbered by inherent limitations such as the inability to fully discern the intricacies of the 3D structure and their reliance on the camera’s perspective.

To surmount these constraints, 3D CNN-based methodologies [15, 31] were introduced, employing 3D voxelized representations of depth images to encapsulate volumetric data. Despite their advancements in 3D hand pose estimation, these techniques demand substantial memory and computational resources, thus curtailing their practical applicability.

In contrast, PointNet-based strategies [5, 8, 13, 16, 24] analyze point clouds, offering a precise depiction of the 3D structure. PointNet, a deep learning framework adept at handling irregular and unstructured point clouds, was initially employed for hand pose estimation in HandPointNet [13]. Subsequent enhancements include the Point-to-Point model [16] and SHPR-Net [5], which refine performance by generating point-wise probability distributions. Notably, SHPR-Net [5] amalgamated HandPointNet with an auxiliary semantic segmentation subnetwork to augment performance. Recently, HandFoldingNet [8] introduced a folding concept to reshape a pre-defined 2D hand skeleton into hand poses, thereby enhancing estimation accuracy. Subsequently, HandR2N2 [6] proposed a folding-based recurrent structure to iteratively optimize the keypoint positions by searching local point regions. However, a notable drawback of point clouds is the computational overhead of querying neighbors from a dense point set for convolution. Therefore, the existing methods commonly use a sparse point cloud, which restricts the performance.

In order to bridge the shortcomings of images and point clouds, IPNet [35] suggests the utilization of multi-modal representations, combining 2D depth images and 3D point clouds. By doing so, the model adeptly extracts dense detail information while effectively capturing 3D spatial features, thereby facilitating accurate 3D hand pose estimation.

Essentially, the above-mentioned methods are all CNN-induced methods but with different input modalities. In contrast, our method introduces a novel adaptive graph transformer architecture for precise local information extraction and dynamic kinematics with the power of the transformer.

2.2 Transformers in Hand Pose Estimation

With the development of the attention mechanism and the introduction of the transformer model [43], various domains within image processing have increasingly embraced its application, such as image classification [9, 28] and object detection [1, 50]. Given the transformer model’s ability to capture non-local features, which proves beneficial for hand pose estimation, numerous studies [19, 21, 23, 25, 26, 45] have endeavored to integrate it into hand pose estimation domain. As a closely related work to ours, Hand-Transformer [21] introduces a non-autoregressive transformer to concurrently estimate keypoint positions. Nevertheless, this approach of independent keypoint estimation overlooks the intrinsic kinematics shared among keypoints. Consequently, HandGCNFormer [45] integrates graph convolutional networks (GCNs) as its core to explicitly model the kinematics among keypoints, thereby enhancing performance. However, its static graph is not able to capture dynamic kinematic correspondence between keypoints under complex occlusion scenarios. In contrast, the proposed HandDAGT is able to dynamically weigh between kinematic correspondence and local detail according to the occlusion situation.

3 Method

The HandDAGT architecture is presented in Figure 2. HandDAGT first samples a super point set with fused features from both image and point cloud, thus it can generate strong keypoint embeddings and informative local patches. The embeddings and patches are further fed as queries and keys to our key component, the adaptive graph transformer. The transformer implements a novel attention mechanism that adaptively weighs queries and keys for the accurate keypoint coordinate estimation.

3.1 Super Points Generation

Following the previous work IPNet [35], HandDAGT accepts the multimodal inputs, which are a hand depth image $\mathbf{D}_{in} \in \mathbb{R}^{H \times W}$ and a set of randomly sampled 3D point coordinates $\mathbf{P}_{in} \in \mathbb{R}^{N \times 3}$. The outputs from HandDAGT are 3D keypoint coordinates $\mathbf{J} \in \mathbb{R}^{J \times 3}$. Since the input point set is irregular and orderless, we exploit the point set convolutional layer [27, 34] as a local 3D encoder to extract a subsampled 3D super point set $\mathbf{P}' \in \mathbb{R}^{N/2 \times 3}$ with corresponding 3D local geometric features $\mathbf{F}_{3d} \in \mathbb{R}^{N/2 \times d_{3d}}$. The depth image is supplied into a ConvNeXt-based autoencoder to generate a 2D local visual feature map $\mathbf{F}_{2d} \in \mathbb{R}^{H/2 \times W/2 \times d_{2d}}$. Subsequently, the 2D local visual features are projected as

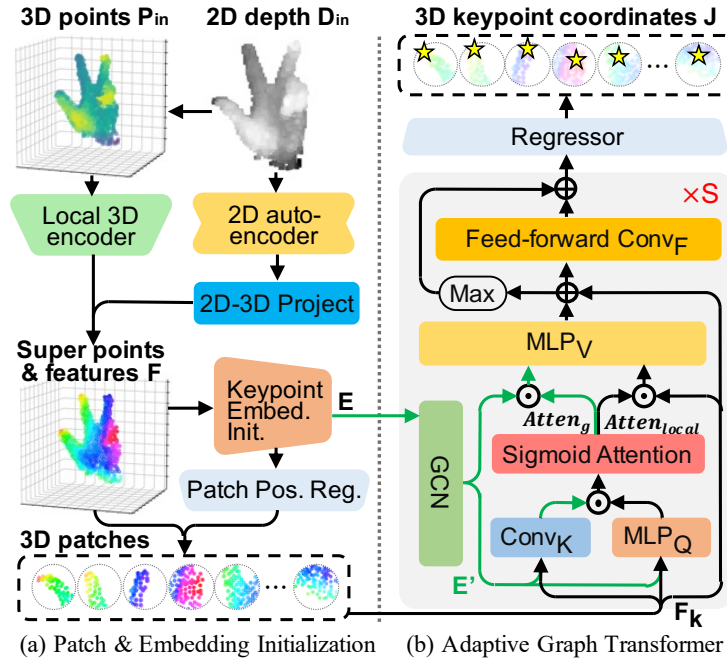


Fig. 2: The HandDAGT architecture. HandDAGT takes a 2D depth image and the sub-sampled point cloud as the input. The PointNet-based local 3D encoder and the 2D auto-encoder extracts local 3D features and local 2D features, respectively. Then, the 2D features are projected into 3D space to fuse with 3D features forming the super point features F . Based on the super points, keypoint embeddings E and 3D patches are extracted as input to the novel adaptive graph transformer to estimate accurate 3D keypoint coordinates by leveraging the dynamic kinematic correspondences and local details. Notably, during the training stage, the local 3D patches are shifted by random noises in order to enforce the model providing robust estimations.

3D point features $\mathbf{F}_{2d-3d} \in \mathbb{R}^{N/2 \times d_{2d}}$ through the 2D-3D projection technique proposed by IPNet [35]. Afterward, the projected features are concatenated to the super points as $\mathbf{F} \in \mathbb{R}^{N/2 \times (d_{2d} + d_{3d})}$. Then the super points are further fed into the initialization module to generate initial keypoint embeddings and patch positions for the transformer.

3.2 Initialization of Keypoint Embedding and Patch Position

In our HandDAGT framework, the keypoint embeddings play a critical role in capturing the kinematic topology between keypoints. To achieve this, we draw upon the concept of joint-wise embeddings initially proposed by HandFoldingNet [8] to obtain the initial keypoint embeddings.

The keypoint embedding initialization module begins by feeding the super points with features \mathbf{F} into two point set convolutional layers to generate 3D

global vectors. The initialization of keypoint embeddings is then sequentially conducted through a three-layer bias-induced layer (BIL) [7]. The global vector is replicated J times and supplied to the BILs. This operation yields embedding for the J individual keypoints. Notably, the BIL furnishes keypoint-wise independent biases, akin to learnable positional embeddings.

After the acquisition of the keypoint embeddings, a linear transformation is applied to project the high-dimensional embeddings into the 3D space as the initial positions for collecting patches. The 3D point patches are acquired by grouping K -nearest super points with their corresponding features around each patch position.

3.3 Adaptive Graph Transformer Decoder

The adaptive transformer decoder is the key component in our proposed framework as it queries the local information for accurate keypoint estimation with the assist of dynamic kinematic topology based on the keypoint embeddings, as illustrated in Figure 2.

The keypoint embeddings are first augmented through a GCN, forming the evolved embeddings,

$$\mathbf{E}' = \text{ReLU}(\mathbf{A}\mathbf{E}\mathbf{W}), \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{d_{hid} \times d_{hid}}$ is the trainable transformation weights and $\mathbf{A} \in \mathbb{R}^{J \times J \times d_{hid}}$ is the learned kinematic correspondence matrix that models connection strength between keypoints in each channel.

Subsequently, each super point in each patch and its corresponding feature are concatenated with the evolved keypoint embedding, forming the input to the novel attention mechanism of the proposed transformer,

$$\mathbf{Q}_k = \text{MLP}_Q([\mathbf{E}', \mathbf{F}_k]), \quad (2)$$

$$\mathbf{K} = \underset{1 \leq k \leq K}{\text{set_conv}}([\mathbf{E}', \mathbf{F}_k]), \quad (3)$$

$$[\mathbf{Att}_k^{\text{local}}, \mathbf{Att}_k^{\text{g}}] = \sigma(\mathbf{Q}_k \odot \mathbf{K}), \quad (4)$$

$$\mathbf{V}_k = \text{MLP}_V([\mathbf{E}' \odot \mathbf{Att}_k^{\text{g}}, \mathbf{F}_k \odot \mathbf{Att}_k^{\text{local}}]), \quad (5)$$

where MLP is the multi-layer perceptron, set_conv is the point set convolutional layer, σ is the sigmoid activation function, $[\cdot, \cdot]$ is the concatenation operation, and \odot is the Hadamard product. We notice that the semantic difference between local points and keypoint embeddings affects the effectiveness of attention, we enhance the queries (local points) \mathbf{Q} with the evolved keypoint embedding while the keys (keypoint embeddings) \mathbf{K} are fused with the local observation. Furthermore, unlike conventional attention that commonly utilizes dot product with softmax, our attention score is obtained by the Hadamard product with a sigmoid as the activation as shown in Equation 4. Moreover,

due to the channel-wise multiplication of the Hadamard product, the proposed attention mechanism can be regarded as inherently employing the multi-head attention with the same number of attention heads as the channel depth. It is worth mentioning that the attention produces two individual scores for the local points and graph-evolved key embeddings. With these two scores, the proposed transformer is able to adaptively focus more on kinematic topology or local information of the keypoints depending on whether the keypoints are occluded or well-represented by the local points.

Following the standard attention block, a residual mechanism is introduced to accumulate the computed value \mathbf{V} and the input local points. Afterward, another point set convolution is applied as the feed-forward module to aggregate the computed value \mathbf{V} to an updated key embedding for position regression of each keypoint,

$$\tilde{\mathbf{E}} = \underset{1 \leq k \leq K}{\text{set_conv}_F}(\mathbf{V}_k + \mathbf{F}_k) + \underset{1 \leq k \leq K}{\text{MAX}_F}(\mathbf{V}_k + \mathbf{F}_k), \quad (6)$$

where MAX_F indicates the max-pooling operation for pursuing the residual accumulation. At the end, a trainable linear transformation is followed to project the updated keypoint embeddings to the 3D keypoint coordinates.

For an improved estimation performance, the proposed transformer is stacked multiple times to progressively optimize the keypoint position. The regressed positions from the previous transformers are used for grouping new patches as the input to the subsequent transformer stacks. The optimal number of stacked transformers will be discussed in Section 4.4.

3.4 Denosing Training and Inference

To improve the accuracy and robustness of the model against noise and occlusion, we propose a novel denoising training strategy. The initial patch positions are firstly disturbed by the introduced random noise. The noise variances during training are controlled by the same linear scheduler as in recent diffusion models [3, 20]. However, in contrast to the diffusion models that introduce noise in both the training and inference phases, our method only introduces noise in the training stage. Thus, the training criterion in our method is formally defined as:

$$\mathcal{L} = \sum L1_{smooth}(\mathbf{D}_{T_1}(\mathbf{J}_0 + \mathcal{N}) - \mathbf{J}^*) + \sum_2^S \sum L1_{smooth}(\mathbf{D}_{T_s}(\mathbf{J}_{s-1}) - \mathbf{J}^*), \quad (7)$$

where \mathcal{N} is the injected noise, S denotes the number of stacked denosing transformers \mathbf{D}_{T_s} , and \mathbf{J}_s is the output from the s -th denosing transformer $\mathbf{D}_{T_s}(\mathbf{J}_{s-1})$. Note that, the noise is only injected in the input of the first denosing transformer. Following the previous works [8, 38], we adopt the smooth L1 loss because of its less sensitivity to the outliers. The smooth L1 loss is defined as:

$$L1_{smooth}(\mathbf{x}) = \begin{cases} 50\mathbf{x}^2, & |\mathbf{x}| < 0.01 \\ |\mathbf{x}| - 0.005, & \text{otherwise} \end{cases}. \quad (8)$$

4 Experiments

4.1 Experiment Settings

In our experiments, we utilized an NVIDIA A100 GPU with PyTorch for training and evaluating the model. We employed the AdamW optimizer [29] with parameters set to $\beta_1 = 0.5$, $\beta_2 = 0.999$, and a learning rate $\alpha = 0.001$. The input images were resized to 128, and the number of input points to the network was sampled to 1,024. We set the batch size to 32. To mitigate overfitting, we employed online data augmentation techniques, including random rotation (in the range of $[-180.0, 180.0]$ degrees), 3D scaling (within the range of $[0.9, 1.1]$), and 3D translation (within the range of $[-10, 10]$ mm). For training on the hand-object datasets, we conducted training for 30 epochs with a learning rate decay of 0.1 after every 10 epochs. In the case of training on the single-hand dataset, we extended the training duration to 50 epochs with a similar learning rate decay strategy, applying the decay after 30 epochs.

4.2 Datasets and Evaluation Metrics

Single-hand Dataset. **ICVL** [40] comprises 22,000 training depth frames and 1,600 testing depth frames. Each frame includes annotations for $J = 16$ keypoints, including one keypoint for the palm and three keypoints for each finger. **NYU** [41] provides depth images of single hands captured from three distinct viewpoints using the PrimeSense 3D sensor. Each viewpoint comprises 72,000 frames for training and 8,000 frames for testing. In line with recent studies [8, 13, 16], a single view is utilized for both training and testing purposes. Furthermore, for evaluation, 14 keypoints out of a total of 36 annotated keypoints are selected, maintaining consistency with prior works.

Hand-object Dataset. **DexYCB** [2] is a recently released hand-object dataset, comprising 582,000 image frames annotated with 21 keypoints. It incorporates data from 10 distinct subjects interacting with 20 YCB objects across 8 camera views. Notably, the dataset defined four official split protocols: S0 - seen subjects, camera views, grasped objects; S1 - unseen subjects; S2 - unseen camera views; S3 - unseen grasped objects. **HO3D** [18] is another challenging dataset renowned for its precise hand-object pose during interactions. The widely utilized version, **HO3D_v2**, comprises 66,034 training images and 11,524 testing images gathered from 10 subjects engaging with 10 distinct objects. Evaluation for **HO3D_v2** is typically conducted through an online submission platform dedicated to this dataset.

Evaluation metrics. We employ two commonly used metrics, the mean keypoint error, and the success rate, to evaluate the performance of hand pose estimation. The mean keypoint error quantifies the average Euclidean distance between the estimated keypoint positions and ground-truth ones for all the keypoint over the entire testing set. The success rate indicates the percentage of good frames with a mean keypoint error of less than a given distance threshold.

Table 1: Comparison of the proposed method with previous state-of-the-art methods on the single-hand ICVL and NYU datasets. Input indicates the input type of 2D depth image (D), 3D voxels (V), or 3D point cloud (P).

Method	Mean keypoint error (mm)		Input
	ICVL	NYU	
Ren-9x6x6 [17]	7.31	12.69	D
DeepPrior++ [32]	8.1	12.24	D
Pose-Ren [4]	6.79	11.81	D
DenseReg [44]	7.3	10.2	D
CrossInfoNet [10]	6.73	10.08	D
JGR-P2O [11]	6.02	8.29	D
SSRN [37]	6.01	7.37	D
PHG [36]	5.97	7.39	D
HandPointNet [13]	6.94	10.54	P
Hand-Transformer [21]	6.47	9.80	P
Point-to-Point [16]	6.3	9.10	P
V2V [31]	6.28	8.42	V
HandFolding [8]	5.95	8.58	P
HandR2N2 [6]	5.70	7.27	P
IPNet [35]	5.76	7.17	D+P
HandDAGT (Ours)	5.66	7.12	D+P

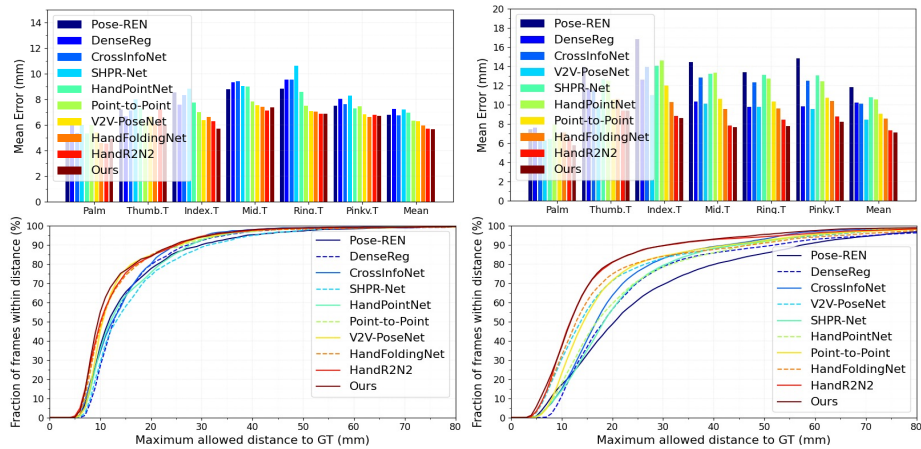


Fig. 3: Comparison with the state-of-the-art methods using the ICVL (left) and NYU (right) dataset. The per keypoint error (top) and success rate (bottom) are shown in this figure.

4.3 Comparison with State-of-the-Art Methods

Single Hand. We compare HandDAGT on the ICVL and NYU dataset with other state-of-the-art methods, including methods with 2D depth images as input: DeepPrior++ [32], Ren-9x6x6 [17], Pose-Ren [4], DenseReg [44], CrossInfoNet [10], JGR-P2O [11], spatial-aware stacked regression network (SSRN)

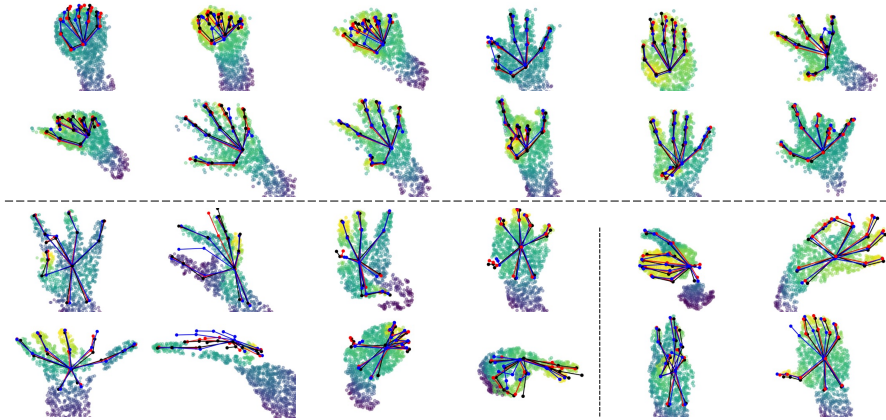


Fig. 4: Qualitative results of HandDAGT on the ICVL (top) and NYU (bottom) datasets. Hand-depth images are transformed into 3D points to clearly illustrate occlusions. Ground truth keypoints are represented in black, results from the comparative HandR2N2 [6] model are shown in blue, and the estimated keypoint coordinates of our model are depicted in red. The bottom figures showcase self-occluded/truncated cases (left) and well-performed cases without occlusion (right) in the NYU dataset.

[37] and pose-guided hierarchical graph network (PHG) [36], methods with 3D point cloud or voxels as input: HandPointNet [13], Point-to-Point [16], Hand-Transformer [21], V2V [31], HandFoldingNet [8] and HandR2N2 [6], and method with multi-modal input, IPNet [35].

Table 1 summarizes the performance in terms of mean keypoint error on two single-hand datasets. The results highlight that HandDAGT achieves a new state-of-the-art record with the lowest mean distance errors of 5.66 mm and 7.12 mm on the challenging ICVL and NYU datasets, respectively. Furthermore, the outcomes demonstrate that HandDAGT significantly outperforms other 2D image-based methods by substantial margins. This superiority is attributed to HandDAGT’s direct processing in the 3D space, which circumvents the highly non-linear mapping problem associated with estimating from 2D images. The qualitative results illustrated in Figure 4 also show the evidence of the HandDAGT’s superiority in both occluded and non-occluded cases. Figure 3 represents that our method significantly outperforms other methods in terms of success rate when the error threshold is lower than 13 and 8 mm on the ICVL and NYU datasets, respectively.

Hand-Object. We compare HandDAGT on the hand-object dataset DexYCB with other state-of-the-art method on the official dataset split protocols, including A2J [47], Spurr et al. [39], METRO [25], Tse et al. [42], HandOccNet [33] and IPNet [35]. As shown in Table 2, HandDAGT significantly outperforms previous SOTA methods on all four protocols. Qualitative results visualized in Figure 5 also reveal that HandDAGT can estimate accurate poses from hand-object interaction scenarios with various occlusions. We also compared HandDAGT




Fig. 5: Qualitative results of HandDAGT on the DexYCB dataset including different grabbing poses (top), interacting objects (2nd row), object occlusions (3rd row), and self-occlusions (bottom). Hand-depth images (left) are transformed into 3D points (right) in order to clearly present occlusions as shown in the figure. Ground truth is shown in black and the estimated keypoint coordinates of our model are shown in colors.

Table 2: Comparison of the proposed method with previous state-of-the-art methods on the DexYCB datasets.

Method	Mean keypoint error (mm)					Input
	S0	S1	S2	S3	AVG	
A2J [47]	23.93	25.57	27.65	24.92	25.52	D
Spurr et al. [39]	17.34	22.26	25.49	18.44	18.44	RGB
METRO [25]	15.24	-	-	-	-	RGB
Tse et al. [42]	16.05	21.22	27.01	17.93	20.55	RGB
HandOcc [33]	14.04	-	-	-	-	RGB
IPNet [35]	8.03	9.01	8.60	7.80	8.36	D+P
HandDAGT (Ours)	7.72	8.68	8.22	7.52	8.03	D+P

on another the hand-object dataset HO3D with state-of-the-art method Hybrik [22], ArtiBoost [48], HandOccNet [33], HandVoxNet++ [30] and IPNet [35]. The comparison results are presented in Table 3. The results demonstrate that our HandDAGT can provide the comparable performance with the same mean distance error compared with the state-of-the-art model IPNet.

4.4 Ablation Study

Effectiveness of the proposed attention mechanism. To assess the effectiveness and necessity of the proposed attention mechanism, we conducted a series of ablation experiments as follows:

Table 3: Comparison of the proposed method with previous state-of-the-art methods on the HO3D datasets.

Method	Mean keypoint error (mm)	Input
Hybrik [22]	2.89	RGB
ArtiBoost [48]	2.53	RGB
HandOccNet [33]	2.49	RGB
HandVoxNet++ [30]	2.46	V
IPNet [35]	1.81	D+P
HandDAGT (Ours)	1.81	D+P

Table 4: Ablations of different configurations of the transformer. All the ablation models are trained and tested on the NYU dataset. ‘C’ denotes the conventional attention mechanism. ‘S’ denotes the proposed sigmoid attention.

ID	Configurations				Mean keypoint error (mm)
	Attention type	Enhanced \mathbf{Q}, \mathbf{K}	$\mathbf{Att}_k^{\text{local}}$	$\mathbf{Att}_k^{\text{g}}$	
1)	×	×	×	×	7.43
2)	C	×	✓	✓	10.66
3)	C	✓	✓	✓	10.38
4)	S	×	✓	✓	7.30
5)	S	✓	✓	×	7.24
6)	S	✓	×	✓	7.21
7)	S	✓	✓	✓	7.12

1) No attention mechanism: Only the feed-forward module is retained for pose estimation. 2) Conventional attention mechanism: The conventional attention mechanism that uses dot product and softmax is implemented. 3) No sigmoid attention: Upon our attention mechanism, the sigmoid attention (Equation 4) is replaced as the dot product and softmax as in the conventional attention. 4) No enhanced \mathbf{Q}, \mathbf{K} : Upon our attention mechanism, the enhancement for query (Equation 2) and key (Equation 3) are removed. 5) Only sigmoid attention on local points $\mathbf{Att}_k^{\text{local}}$: Attention mechanism solely for local points is activated. 6) Only sigmoid attention on kinematic topology $\mathbf{Att}_k^{\text{g}}$: Attention mechanism concentrating exclusively on kinematic topology is activated. 7) Full implementation: Both sigmoid attentions on local points and kinematic topology are activated. It’s worth noting that the denoising training strategy is applied for all ablations, and the number of stacked transformers is set to 3 for consistency and comparison across experiments.

Table 4 presents the experimental results of the ablations, which illustrate the contributions of the proposed components to the final performance. The comparison between ablations 1) and 2) highlights that the conventional attention mechanism alone cannot effectively process keypoint coordinate estimation and even deteriorate performance. Conversely, the proposed sigmoid attention mechanism (Equation 4) leads to a drastic improvement in performance, as evidenced




Fig. 6: Qualitative results of HandDAGT ablations on the NYU dataset. Ground truth is shown in black and the estimated keypoint coordinates of ablations 1) no attention, 5) attention on local points $\text{Att}_k^{\text{local}}$, 6) attention on kinematic topology Att_k^g and 7) the full implementation are shown in blue, orange, green and red, respectively.

Table 5: Ablation experimental result. All the ablation models are trained and tested on the NYU dataset.

(a) Denoising training.		(b) Number of transformer.		(c) Size of input patch.	
Denosing training	Mean keypoint error (mm)	Number of transformer	Mean keypoint error (mm)	Patch size (pts)	Mean keypoint error (mm)
×	7.20	1	8.02	32	7.20
✓	7.12	2	7.15	64	7.12
		3	7.12	128	7.16

by the comparison between ablations 1) and 4). Furthermore, the results from ablations 3) and 4) suggest that the enhanced \mathbf{Q} and \mathbf{K} contribute significantly to the model’s performance, as well as the conventional attention mechanism. Ablations 5) and 6) demonstrate that both $\text{Att}_k^{\text{local}}$ and Att_k^g positively impact the model’s performance, indicating that attention mechanisms focused on local points and kinematic topology are beneficial additions to the model architecture. We also visualize the qualitative results of the ablations in Figure 6. The qualitative results indicate that attention solely on local points can yield better performance in non-occluded cases, as illustrated in the left of Figure 6. This is attributed to the sufficient local details. Conversely, attention solely on kinematic topology demonstrates strong performance in occluded cases, as depicted in the right of Figure 6. Our full implementation surpasses all ablations by leveraging both mechanisms, thereby achieving superior performance.

Effectiveness of the denoising training strategy. To evaluate the effectiveness of the proposed denoising training strategy, we train a model without noise disturbing during training. As reported in Table 5a, the performance is boosted with 0.08 mm of error reduction. Notably, this performance improvement does not require any additional computations, since these two models share the equivalent inference process but vary with different training strategies.

Number of stacked transformers. Intuitively, deeper stacks of transformers have the potential to enhance performance by allowing the model to capture more intricate features. To explore this, we incrementally increased the number of stacked transformers to search for the optimal depth of the transformer stack. As outlined in Table 5b, we systematically varied the depth of the stack to evaluate its impact on model performance. Based on the results, it is evident that increasing the depth of the transformer stack generally leads to performance

improvements. However, the model with a 3-stack transformer exhibits only marginal enhancements. Therefore, we conclude that the optimal stack number for our model is 3, and further increasing the depth of the transformer stack does not yield significant performance gains.

Size of input patches. Indeed, the size of patches significantly influences the receptive field for the transformer, ultimately impacting the model’s ability to capture detailed information. To investigate this effect, we conducted a series of ablation experiments wherein the models were trained using different patch sizes for estimation. Since the patches consist of super points, thus the size of the patch is the number of super points contained in a patch. As depicted in Table 5c, the patch containing 64 super points exhibited optimal performance, whereas the experiment with patch sizes larger than 64 points showed degraded performance. This observation suggests that the receptive field of the transformer must be adequately sized to gather effective information. However, introducing a redundant patch size could lead to the inclusion of unnecessary information, thus disrupting the transformer’s operation.

5 Conclusion

In this paper, we introduced HandDAGT, a novel adaptive graph transformer architecture designed to adaptively utilize local information and kinematic topology based on the occlusion circumstances. To improve the robustness of the proposed network, a novel denoising training strategy is introduced. Experimental results showcased that our network outperforms previous state-of-the-art methods on four challenging datasets. Moreover, extensive experiments validated the effectiveness of the components proposed in this paper. As a limitation, HandDAGT currently can not process the scenario with interacting hands and requires redundant computation due to the 2D feature extraction. Addressing these limitations could involve extending our research to include bidirectional learning and developing a lightweight 2D model. These avenues remain open for future research and represent potential areas for enhancing the capabilities and efficiency of HandDAGT.

Acknowledgement

This work was supported by the Ministry of Science and ICT (MSIT) of Korea under the National Research Foundation (NRF) grant (2022R1A4A3032913) and the Institute of Information and Communication Technology Planning Evaluation (IITP) grants (IITP-2019-0-00421, IITP-2023-2020-0-01821, IITP-2021-0-02052, IITP-2021-0-02068), and by the Technology Innovation Program (RS-2023-00235718) funded by the Ministry of Trade, Industry & Energy (1415187474). Wencan Cheng was partly supported by the China Scholarship Council (CSC).

References

1. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
2. Chao, Y.W., Yang, W., Xiang, Y., Molchanov, P., Handa, A., Tremblay, J., Narang, Y.S., Van Wyk, K., Iqbal, U., Birchfield, S., et al.: Dexycb: A benchmark for capturing hand grasping of objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9044–9053 (2021)
3. Chen, S., Sun, P., Song, Y., Luo, P.: Diffusionnet: Diffusion model for object detection. arXiv preprint arXiv:2211.09788 (2022)
4. Chen, X., Wang, G., Guo, H., Zhang, C.: Pose guided structured region ensemble network for cascaded hand pose estimation. *Neurocomputing* **395**, 138–149 (2020)
5. Chen, X., Wang, G., Zhang, C., Kim, T.K., Ji, X.: Shpr-net: Deep semantic hand pose regression from point clouds. *IEEE Access* **6**, 43425–43439 (2018)
6. Cheng, W., Ko, J.H.: Handr2n2: Iterative 3d hand pose estimation using a residual recurrent neural network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 20904–20913 (2023)
7. Cheng, W., Lee, S.: Point auto-encoder and its application to 2d-3d transformation. In: International Symposium on Visual Computing. pp. 66–78. Springer (2019)
8. Cheng, W., Park, J.H., Ko, J.H.: Handfoldingnet: A 3d hand pose estimation network using multiscale-feature guided folding of a 2d hand skeleton. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11260–11269 (2021)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
10. Du, K., Lin, X., Sun, Y., Ma, X.: Crossinfonet: Multi-task information sharing based hand pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9896–9905 (2019)
11. Fang, L., Liu, X., Liu, L., Xu, H., Kang, W.: Jgr-p2o: Joint graph reasoning based pixel-to-offset prediction network for 3d hand pose estimation from a single depth image. In: European Conference on Computer Vision. pp. 120–137. Springer (2020)
12. Gao, Y., Liu, X., Li, J., Fang, Z., Jiang, X., Huq, K.M.S.: Lft-net: Local feature transformer network for point clouds analysis. *IEEE transactions on intelligent transportation systems* **24**(2), 2158–2168 (2022)
13. Ge, L., Cai, Y., Weng, J., Yuan, J.: Hand pointnet: 3d hand pose estimation using point sets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8417–8426 (2018)
14. Ge, L., Liang, H., Yuan, J., Thalmann, D.: Robust 3d hand pose estimation in single depth images: from single-view cnn to multi-view cnns. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3593–3601 (2016)
15. Ge, L., Liang, H., Yuan, J., Thalmann, D.: 3d convolutional neural networks for efficient and robust hand pose estimation from single depth images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1991–2000 (2017)
16. Ge, L., Ren, Z., Yuan, J.: Point-to-point regression pointnet for 3d hand pose estimation. In: Proceedings of the European conference on computer vision (ECCV). pp. 475–491 (2018)

17. Guo, H., Wang, G., Chen, X., Zhang, C., Qiao, F., Yang, H.: Region ensemble network: Improving convolutional network for hand pose estimation. In: 2017 IEEE International Conference on Image Processing (ICIP). pp. 4512–4516. IEEE (2017)
18. Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3d annotation of hand and object poses. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3196–3206 (2020)
19. Hampali, S., Sarkar, S.D., Rad, M., Lepetit, V.: Keypoint transformer: Solving joint identification in challenging hands and object interactions for accurate 3d pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11090–11100 (2022)
20. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
21. Huang, L., Tan, J., Liu, J., Yuan, J.: Hand-transformer: Non-autoregressive structured modeling for 3d hand pose estimation. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV* 16. pp. 17–33. Springer (2020)
22. Li, J., Xu, C., Chen, Z., Bian, S., Yang, L., Lu, C.: Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3383–3393 (2021)
23. Li, M., An, L., Zhang, H., Wu, L., Chen, F., Yu, T., Liu, Y.: Interacting attention graph for single image two-hand reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2761–2770 (2022)
24. Li, S., Lee, D.: Point-to-pose voting based hand pose estimation using residual permutation equivariant layer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11927–11936 (2019)
25. Lin, K., Wang, L., Liu, Z.: End-to-end human pose and mesh reconstruction with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1954–1963 (2021)
26. Lin, K., Wang, L., Liu, Z.: Mesh graphormer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12939–12948 (2021)
27. Liu, X., Qi, C.R., Guibas, L.J.: Flownet3d: Learning scene flow in 3d point clouds. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 529–537 (2019)
28. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
29. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
30. Malik, J., Shimada, S., Elhayek, A., Ali, S.A., Theobalt, C., Golyanik, V., Stricker, D.: Handvoxnet++: 3d hand shape and pose estimation using voxel-based neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(12), 8962–8974 (2021)
31. Moon, G., Yong Chang, J., Mu Lee, K.: V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5079–5088 (2018)
32. Oberweger, M., Lepetit, V.: Deeprior++: Improving fast and accurate 3d hand pose estimation. In: Proceedings of the IEEE international conference on computer vision Workshops. pp. 585–594 (2017)

33. Park, J., Oh, Y., Moon, G., Choi, H., Lee, K.M.: Handocnet: Occlusion-robust 3d hand mesh estimation network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1496–1505 (2022)
34. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Advances in neural information processing systems. pp. 5099–5108 (2017)
35. Ren, P., Chen, Y., Hao, J., Sun, H., Qi, Q., Wang, J., Liao, J.: Two heads are better than one: image-point cloud network for depth-based 3d hand pose estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 2163–2171 (2023)
36. Ren, P., Sun, H., Hao, J., Qi, Q., Wang, J., Liao, J.: Pose-guided hierarchical graph reasoning for 3-d hand pose estimation from a single depth image. *IEEE Transactions on Cybernetics* (2021)
37. Ren, P., Sun, H., Huang, W., Hao, J., Cheng, D., Qi, Q., Wang, J., Liao, J.: Spatial-aware stacked regression network for real-time 3d hand pose estimation. *Neurocomputing* **437**, 42–57 (2021)
38. Ren, P., Sun, H., Qi, Q., Wang, J., Huang, W.: Srn: Stacked regression network for real-time 3d hand pose estimation. In: *BMVC*. p. 112 (2019)
39. Spurr, A., Iqbal, U., Molchanov, P., Hilliges, O., Kautz, J.: Weakly supervised 3d hand pose estimation via biomechanical constraints. In: *European conference on computer vision*. pp. 211–228. Springer (2020)
40. Tang, D., Jin Chang, H., Tejani, A., Kim, T.K.: Latent regression forest: Structured estimation of 3d articulated hand posture. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3786–3793 (2014)
41. Tompson, J., Stein, M., Lecun, Y., Perlin, K.: Real-time continuous pose recovery of human hands using convolutional networks. *ACM Transactions on Graphics (ToG)* **33**(5), 1–10 (2014)
42. Tse, T.H.E., Kim, K.I., Leonardis, A., Chang, H.J.: Collaborative learning for hand and object reconstruction with attention-guided graph convolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1664–1674 (2022)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
44. Wan, C., Probst, T., Van Gool, L., Yao, A.: Dense 3d regression for hand pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5147–5156 (2018)
45. Wang, Y., Chen, L., Li, J., Zhang, X.: Handgcnformer: A novel topology-aware transformer network for 3d hand pose estimation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5675–5684 (2023)
46. Wu, L., Liu, X., Liu, Q.: Centroid transformers: Learning to abstract with attention. *arXiv preprint arXiv:2102.08606* (2021)
47. Xiong, F., Zhang, B., Xiao, Y., Cao, Z., Yu, T., Zhou, J.T., Yuan, J.: A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 793–802 (2019)
48. Yang, L., Li, K., Zhan, X., Lv, J., Xu, W., Li, J., Lu, C.: Artiboost: Boosting articulated 3d hand-object pose estimation via online exploration and synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2750–2760 (2022)

49. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 16259–16268 (2021)
50. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)